arXiv:1612.03052v1 [cs.CV] 9 Dec 2016

ActionFlowNet: Learning Motion Representation for Action Recognition

Joe Yue-Hei Ng[1] Jonghyun Choi[2] Jan Neumann[3] Larry S. Davis[1]
[1]University of Maryland, College Park [2]Allen Institute for AI [3]Comcast Labs, DC
{yhng,lsd}@umiacs.umd.edu jonghyunc@allenai.org jan_neumann@cable.comcast.com

Abstract

Even with the recent advances in convolutional neural networks (CNN) in various visual recognition tasks, the state-of-the-art action recognition system still relies on hand crafted motion feature such as optical flow to achieve the best performance. We propose a multitask learning model ActionFlowNet to train a single stream network directly from raw pixels to jointly estimate optical flow while recognizing actions with convolutional neural networks, capturing both appearance and motion in a single model. We additionally provide insights to how the quality of the learned optical flow affects the action recognition. Our model not only significantly improves action recognition accuracy by a large margin (17%) *compared to state-of-the-art CNN-based action recognition models trained without external large scale data and additional optical flow input, but also produces the optical flow as a side product.*

1. Introduction

Recognizing actions in videos is one of the most active research areas with numerous applications including visual surveillance, human-machine interfaces, video retrieval and various visual contents analysis [2]. An action is a sequence of movements of one or more objects. Thus, object motion is important for learning a discriminative action classifier. There are a number of ways to represent motion for recognizing actions: trajectories of local patches [25], hand-crafted spatio-temporal features [26], learned spatio-temporal feature [8] and deep features learned from optical flow [17], which is a pixel-wise short term motion representation.

Two-stream convolutional neural networks, which separately learn appearance and motion by two convolutional networks on static images and optical flow respectively, show impressive results on action recognition [17]. The separation, however, fails to learn the interaction between the motion and the appearance of objects and introduces additional complexity of computing the flow to the classification pipeline. Instead, we propose a model to learn both motion features and appearance directly from raw pixels without hand-crafted flow input.

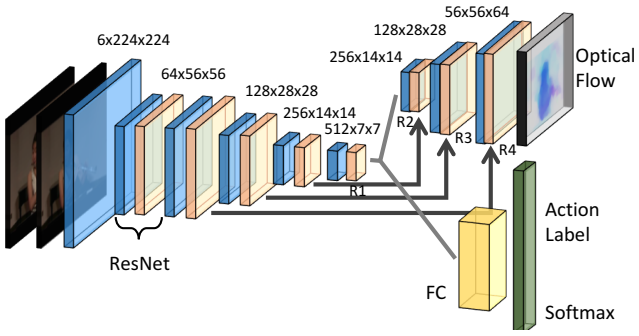

Figure 1: ActionFlowNet for jointly estaimting optical flow and recognizing actions. Yellow and blue blocks represent ResNet modules, where blue blocks represents strided convolution.

Given large scale datasets such as Sports-1M [8], one could train a high capacity classifier to learn complex motion signatures for action recognition by extending image based CNN architectures with 3D convolutions for video action recongition [8, 21]. While classification loss is an excellent generic appearance learner for image classification, it is not necessarily the most effective supervision for learning motion features for action recognition. As shown in [21], even with large amounts of labeled appearance data, the model is still inferior to models trained on optical flow from small datasets on action recognition, which suggests that an effective representation learning method may be more important than a large amount of training data. Motivated by this, we explore how to effectively learn motion representations for action recognition without large amounts of external labeled data.

Encouraged by the success on estimating optical flow with convolutional neural networks [5], we train a single stream feed-forward convolutional neural network - *ActionFlowNet* for jointly recognizing actions and estimating optical flow. Specifically, we formulate the learning problem as multitask learning, which enables the network

to learn both appearance and motion in a single network from raw pixels. The proposed architecture is illustrated in Figure 1.

Our experiments and analyses show that our model successfully learns motion features for action recognition and provide insights on how the learned optical flow quality affects action classification. We demonstrate the effectiveness of our learned motion representation on two standard action recognition benchmarks - UCF101 and HMDB51. Without providing external training data or fine-tuning from already well-trained model with millions of samples, we show that jointly learning action and optical flow significantly boosts action recognition accuracy compared to state-of-the-art representation learning methods trained without external labeled data.

2. Related Work

Over the past few years, action recognition accuracy is greatly improved by learned features and various learning models utilizing deep convolutional and recurrent neural networks. Please refer to a comprehensive survey of recent advances in action recognition by Cheng *et al*. [2].

We utilize the deep model for action recognition and review recent related work. Simonyan and Zisserman proposed a two stream network architecture to recognize action using both appearance and motions separately [17]. A number of follow up work is proposed based on two-stream networks and further improved action recognition accuracies [4, 28, 27]. Our work is motivated by their success of incorporating optical flow for motion feature learning, but we focus on learning from raw pixels instead of handcrafted representations.

Pre-training the network with a large dataset helps to learn appearance signatures for action recognition. Karpathy *et al*. proposed a "Slow Fusion" network which uses 3D convolutions for large scale video classification [8]. Tran *et al*. trained a 3D convolutional neural network (C3D) with a large amount of data and showed the learned features are generic for different tasks [21]. Without training the model with large dataset, our model performs comparably to models trained with large datasets.

Ng *et al*. developed the network architecture using feature pooling and the long short term memory (LSTM) module and trained with very long video clips [14]. In contrast, we focus on learning short term motion features with only two frames. The long term information aggregation techniques are complementary to ours and could further improve accuracy when integrated.

Our work is based on the recent advances of CNN architecture. He *et al*. proposed a layer skipping architecture for ease of training a high capacity deep model, called residual network or ResNet [6]. We use the ResNet module for ease of training the high capacity deep model in this work.

Optical flow encodes motion between frames and is highly related to action recognition. EpicFlow [16] is one of the best performing optical flow estimation algorithms. It densifies the sparse matching of correspondences by edge-preserving interpolation followed by a variational energy minimization to obtain fine-grained flow estimation. The optical flow estimation, however, also has been noticeably improved due to the success of deep convolutional neural network. Our model is motivated by the success of FlowNet, but emphasizes on improving action recognition.

Leveraging videos as source for unsupervised learning has been suggested recently. Different surrogate tasks have been proposed to learn visual representation from videos without any label. Wang *et al*. trained a network to learn visual similarity for patches obtained from visual tracking in videos [29]. Misra *et al*. trained a network to differentiate the temporal order of different frames from a video [13]. Our work is in similar spirit that to use video as an additional source for learning visual representation. However, different from previous work which focused on learning visual representation for a single image, we learn motion representation for videos which models more than a single frame. Vondrick *et al*. proposed to use a Generatie Adversarial Network to learn a generative model for video [24]. We focus on learning motion representation but not video generation.

Concurrent to our work, Diba *et al*. trained a two stream network with flow prediction [3]. They based their network on C3D with a two-stream architecture. Our work employs a single stream network to learn both appearance and motion. We focus on learning motion representation from scratch and provide more analysis to learn flow representation for action recognition.

3. Approach

We propose a single end-to-end model to learn both motions and action classes simultaneously. Our primary goal is to improve action classification accuracy with the help of motion information; we use optical flow as a motion signature. Unlike previous methods that utilize externally computed optical flow as the input to their models, we only use the video frames for input and simultaneously learn the flow and class labels.

Specifically, we use two consecutive frames for the input for three reasons. First, when there are multiple frames in the input, it is not trivial to analyze whether the performance improvement comes from motion modeling or aggregating long term appearance information. Thus for better analysis, it is desirably to use the two frame input. Second, training two-frame models is computationally much more efficient than multi-frame models, which take N video frames and output $N-1$ optical flow images. Third, existing large scale optical flow datasets, such as the FlyingChairs dataset [5],

provide ground-truth flow on only two consecutive frames, which makes training of optical flow estimation models for many frames difficult.

3.1. Learning Flows with Residual Network

Fischer et al. recently proposed FlowNet [5] that is based on convolutional neural networks to estimate high quality optical flow. In addition, He et al. introduced a residual networks (ResNet) to train a high capacity (deeper) convolutional neural network model by skipping convolutional layers by shortcut connections. Taking advantage of ResNet for the flow estimation, we propose an architecture that is similar to FlowNet-S in structure but with ResNet blocks.

In addition to the benefit of easy training, ResNet is fully convolutional which is easily applied to pixel-wise prediction of optical flow, unlike many architectures with fully connected layers including AlexNet [10] and VGG-16 [18]. In contrast to other classification architecture like AlexNet and VGG-16, which contains multiple max pooling layers that may harm optical flow estimation, the ResNet architecture only contains one pooling layer right after conv1. We believe the reduced number of pooling layers makes ResNet more suitable for optical flow estimation where spatial details need to be preserved. Specifically, we use an 18 layers ResNet, which is computationally efficient with good classification performance [6]. Like FlowNet-S, we concatenate two consecutive frames to produce a $6(\text{ch}) \times 224(\text{w}) \times 224(\text{h})$ input. Skip connections from encoder to decoder are retained to obtain higher resolution information in the decoder.

Following [5], the network is optimized over the end-point error (EPE), which is the sum of L_2 distance between the ground truth optical flow and the obtained flow over all pixels. At the decoder, there are four outputs with different resolutions. The total optical flow loss is the weighted sum of end-point error at multiple resolutions per the following equation:

$$\sum_{r=1}^{4} \alpha_r \sum_{p} \|\mathbf{o}_{j,i,p}^{(r)} - \widehat{\mathbf{o}_{j,i,p}^{(r)}}\|_2, \qquad (1)$$

where $\mathbf{o}_{j,i,p}^{(r)}$ is 2-dimensional optical flow vector of the r^{th} layer output of the i^{th} and the $(i+1)^{th}$ frame in the j^{th} video at pixel p. α_r is the weighting coefficient of the r^{th} optical flow output.

3.2. Learning Action Classes with the Flow

We propose two strategies to transfer the motion information to video action recognition task. We first propose an architecture to classify actions on top of the optical flow estimation network, which we call the *Stacked Model*. Then, we propose an architecture to classify the actions and estimate the optical flow in a multitask learning framework, which we call the *ActionFlowNet*.

Knowledge Transfer by Finetuning. Finetuning a pre-trained network is a common practice to transfer knowledge from different datasets and tasks. Unlike previous work, where knowledge transfer has been accompolished between very similar tasks (image classification and detection or semantic segmentation), knowledge transfer in our model is challenging since the goals of pixel-wise optical flow and action classification are not obviously compatible. We transfer the learned motion by initializing the classification network using a network trained for optical flow estimation; we refer to the pre-trained optical flow estimation network as *FlowNet*. Since the network was trained to predict optical flow, it should encode motion information in intermediate levels which support action classification. However, transfering network is known to have the problem of catastrophic forgetting. Specifically, when training the network for action recognition, the originally initialized flow information could be destroyed when the network adapts the appearance information. We prevent catastrophic forgetting by using the multitask learning framework.

3.2.1 Stacked Model

A straightforward way to use the trained parameters from FlowNet is to take the output of FlowNet and learn a CNN on top of the output, as shown in Figure 2. This is reminiscence of the temporal stream in [17] which learns a CNN on precomputed optical flow. If the learned optical flow has high quality, it should give similar performance to learning a network on optical flow.

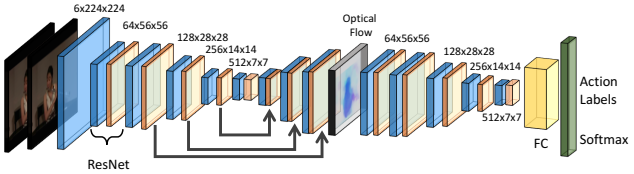

Figure 2: Network structure of the 'Stacked Model'.

Given two frames as input (i^{th} and $(i+1)^{th}$ frame of j^{th} video), the loss function is defined for the first frame of the pair, i^{th} frame, by a standard cross entropy classification loss:

$$\text{Stacked-Loss}_{j,i} = -\mathbb{1}(y_j = \widehat{y_j}) \log p(\widehat{y_j}), \qquad (2)$$

where $\mathbb{1}(\cdot)$ is the indicator function and y_j is the action label of j^{th} video. $\widehat{y_j}$ represents the predicted probability of class $\widehat{y_j}$.

Since the output of FlowNet has 4 times lower resolution than the original image, we remove the first two layers of the CNN (conv1 and pool1) and stack the network on top of it. We also tried to upsample the flow to the original resolution and use the original architecture including conv1

and pool1, but this produces slightly worse results and is computationally more expensive.

The stacked model introduces about 2x number of parameters compared to the original ResNet, and is also 2x more expensive for inference. It learns motion features by explicitly including optical flow as an intermediate representation, but cannot model appearance and motion simultaneously, similar to learning a CNN on precomputed optical flow.

3.2.2 ActionFlowNet

To force the model to learn motion features while training for action recognition, we propose a multtask model which simultaneously learns to estimate optical flow and action classification to avoid catastrophic forgetting. With optical flow as supervision, the model can effectively learn motion features while not relying on explicit optical flow computation. The architecture is shown in Figure 1.

The multitask ActionFlowNet architecture is based on FlowNet with additional classsification layers. Similar to [6], classification is performed by average pooling the last convolutional layer in the encoder followed by a linear classsifier. It is worth noting that this is a very general architecture and requires minimal architectural engineering. Thus, it can be trivially extended to learn more tasks jointly to adapt knowledge from different domains.

Just as with the stacked model, the loss function is defined for each frame. For the i^{th} frame in the j^{th} video the loss is defined as a weighted sum of classification loss and optical flow loss:

$$\text{MT-Loss}_{j,i} = \underbrace{-\mathbb{1}(y_j = \widehat{y_j}) \log p(\widehat{y_j})}_{\text{Classification Loss}} + \underbrace{\lambda \sum_{r=1}^{4} \alpha_r \sum_{p} \|\mathbf{o}_{j,i,p}^{(r)} - \widehat{\mathbf{o}_{j,i,p}^{(r)}}\|_2}_{\text{Flow Loss}}, \quad (3)$$

where $\mathbb{1}(\cdot)$ is a indicator function, y_j and $\widehat{y_j}$ are the groundtruth and predicted action labels respectively of the j^{th} video. $\mathbf{o}_{j,i,p}^{(r)}$ is the optical flow of the r^{th} resolution output of the i^{th} and $(i+1)^{th}$ frame in the j^{th} video at pixel p. α_r is a weighting coefficient for the r^{th} optical flow output. λ is a hyper-parameter balancing the classification loss and the flow loss, where optical flow estimation can be seen as a regularizer for the model to learn motion feature for classification.

Although previous work on multitask learning [12] suggests that sharing parameters of two different tasks may hurt performance, this architecture performs well since optical flow is known empirically to improve video action recognition significantly. In addition, our architecture contains multiple skip connections from lower convolutional layers to decoder. This allows higher layers in the encoder to focus on learning more abstract and high level features, without constraining them to remembering all image details for predicting optical flow, which is beneficial for action recognition. This idea is central to Ladder Network [15] which introduced lateral connections to learn denoising functions and significantly improved classification performance.

4. Experiments

4.1. Datasets

We use two publicly available datasets, UCF101 and HMDB51, to evaluate action classification accuracy. The UCF101 dataset contains 13,320 videos with 101 action classes [19]. The HMDB51 contains 6,766 videos with 51 action categories [11]. As the number of training videos in HMDB51 is small, we initialized our models trained on UCF101 and fine-tuned for HMDB51.

The UCF101 and HMDB51 do not have groundtruth optical flow annotation. Similar to [22], we obtain EpicFlow [16] and use it as a psuedo-groundtruth optical flow to train the motion part of the network.

To learn the motion signature better, we also use FlyingChairs dataset [5] as it has groundtruth optical flow since it is a synthesized dataset. The FlyingChairs dataset contains 22,872 image pairs and ground truth flow from synthetically generated chairs on real images. To validate the quality of optical flow obtained by our models, we use the Sintel dataset [1], which provides dense groundtruth optical flow.

4.2. Experimental Setup

Overfitting Prevention. To prevent overfitting in training, we use different data augmentations on different datasets and different tasks, and use dropout [10]. On the FlyingChairs dataset for optical flow estimation, we augment the data using multi-scale cropping, horizontal flipping, translation and rotation following [5]. On the UCF101 dataset for optical flow estimation, we use multi-scale cropping and horizontal flipping, but do not use translation and rotation in order to maintain the original optical flow distribution in the data. On UCF101 dataset for action recognition, we use color jittering [20], multi-scale cropping and horizontal flipping. Dropout is only applied to the output of the average pooling layer before the linear classifier with probability 0.5.

Optimization and Evaluation. The models are trained using Adam [9] for 80,000 iterations with learning rate 3×10^{-4}. For evaluation, we sample 64 random crops from two random consecutive frames from a video and run a forward pass to the network and average the prediction scores.

4.3. Improving Action Recognition

We first evaluate the action recognition accuracy by the various proposed models, described in Section 3.2, on both UCF101 and HMDB51 datasets. The recognition accuracies are summarized in Table 1.

Method	UCF101	HMDB51
Scratch	47.7	23.9
Fine-tune	66.0	29.1
Stacked	69.6	42.4
ActionFlowNet (UCF101)	70.0	42.4
ActionFlowNet (FlCh+UCF101)	**71.0**	**42.6**
ImageNet pretrained ResNet-18	80.7	45.0

Table 1: Action recognition accuracies of our models on UCF101 and HMDB51 datasets (split 1). FlCh denotes FlyingChairs dataset. "ActionFlowNet (UCF101)" denotes its FlowNet part is pretrained on UCF101, and "ActionFlowNet (FlCh+UCF101)" denotes its FlowNet part is pretrained on FlyingChairs dataset. Both ActionFlowNets are then learned on UCF101 dataset for action and flow.

'Scratch' is a ResNet-18 model that is trained from scratch using UCF101 without any extra motion information. 'Fine-tune' is a model that is trained from UCF101 for optical flow and fine-tuned with action classification. It captures motion information by initialized FlowNet but not in the fine-tuning stage. 'Stacked' is a stacked classification model on top of optical flow output depicted in Figure 2. Its underlying FlowNet is trained with UCF101 and is fixed to predict optical flow, so only the CNN classifier on top is learned. 'ActionFlowNet' is the multitask model depicted in Figure 1, which is trained for action recognition and optical flow estimation to learn both motion and appearance. We trained two versions of ActionFlowNet: one with FlowNet pretrained on UCF101 and one with on FlyingChairs dataset.

As shown in the table, all proposed models - 'Fine-tune', 'Stacked' and 'ActionFlowNet' significantly outperform 'Scratch' . This implies that our models can take advantage of the learned motion for action recognition, which is difficult to learn implicitly from action labels.

Both the Stacked model and two ActionFlowNets outperform the finetuning models by a large margin (up to 5.0% in UCF101 and up to 13.5% in HMDB51). As all models are pretrained from the high quality optical flow model, the results show that the knowledge learned from previous task is prone to be forgotten when learning new task without multitask learning. With extra supervision from optical flow estimation, multitask models regularize the action recognition with the effort of learning the motion features.

While the Stacked model performs similarly to ActionFlowNet when trained only on UCF101, ActionFlowNet is much more compact than the Stacked model, containing only approximately half the number of parameters of the Stacked model. When ActionFlowNet is first pretrained with FlyingChairs, which predicts better quality optical flow in EPE, and finetuned with the UCF101, it further improves accuracy by 1%. This implies that our multitask model is capable of transferring general motion information from other datasets to improve recognition accuracy further.

Our ActionFlowNet still performs inferior compared to ResNet pretrained on ImageNet, especially in UCF101 (71.0% vs 80.7%) because of the rich background context appearance in the dataset. When evaluated on HMDB51, where the backgrounds are less discriminative, our ActionFlowNet is only slightly behind the ImageNet pretrained model (42.6% vs 45.0%), indicating that our model learns strong motion features for action recognition.

Compared to the temporal stream models [17], ActionFlowNet runs significantly faster at test time (0.005 second per image in mini-batch) by orders of magnitude in seconds, as optical flow computation (about 0.06 second per image) is not needed at test time.

Comparison to state-of-the-arts. We compare our approach to previous work in Table 2 on UCF101. For fair comparison, we present the results which only uses UCF101 labels for training. Our models significantly outperform previous work that use videos for unsupervised feature learning. Specifically, even with only our fine-tuned model on UCF101, we obtain 13% improvement compared to VGAN and Sequential Verification, indicating the importance of motion in learning video representations. When combined with multitask learning, the performance improves to 70.0%, which gives 17% improvement over previous work. This shows that our proposed framework can effectively learn motion features for video classification, which is not explicitly learned in previous work.

Method	UCF101 Accuracy
ResNet-18 Scratch	47.3
AlexNet Scratch [17]	52.9
Sequential Verification [13]	50.9
VGAN [24]	52.1
FlowNet fine-tuned (ours)	66.0
ActionFlowNet (ours)	**70.0**

Table 2: Results on UCF101 (split 1) from single stream networks with raw pixel input.

4.3.1 Learning Motions for Discriminative Regions

We visualize what is learned from the multitask network by using the method from Zeiler and Fergus [31]. We fill a black square to occlude the frames at different spatial loca-

(a) Image (b) Output flow from ActionFlowNet (c) Appearance Only (d) ActionFlowNet: Motion and Appearance

Figure 3: Visualization of important regions for action recognition. 'Appearance Only' refers to the output of ImageNet pretrained ResNet-18. Our ActionFlowNet discovers the regions where the motions are happening to be important while 'Appearance Only' captures discriminative regions based on the appearance.

tions and compute the relative difference between classification confidence before and after occlusion.

We compare the most discriminative regions discovered by our multitask network with ones by the ImageNet trained ResNet-18, which only models the discriminative appearances without motion. Figure 3 shows example results. The visualization reveals that our model focuses more on motion, while the ImageNet pretrained network relies more on background appearance, which may not directly relate to the action itself. However, when appearance is discriminative - for example the writing on the board in the last example - our model can also focus on appearance, which is not possible for models that learn from optical flow only.

4.3.2 Classes Improved By the Motions

Not all action classes are motion-centric - objects and their contextual (background) appearances provide more discriminative information for some classes [7]. We present per class accuracy for the models trained jointly with motion (ActionFlowNet) and without multitask (fine-tuned FlowNet).

The top 5 classes with the most improvement (amount in parenthesis) using motion information are JumpingJack (+59%), PushUps (+50%), JugglingBalls (+33%), BabyCrawling (+31%) and BenchPress (+29%). The top 5 classes with the most degraded performance (amount in parenthesis) using motion information are MoppingFloor (-21%), SoccerPenality (-17%), Fencing (-15%), PlayingDhol (-12%) and Hammering (-12%). Figure 5 shows one example of improved and degraded from each classes.

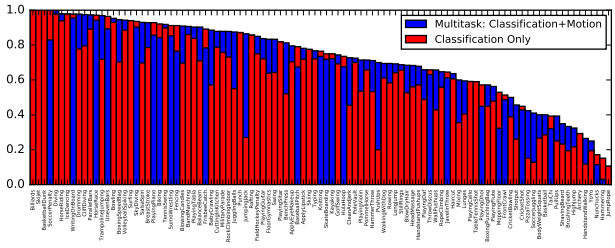

Figure 4: Classwise accuracy improvement by ActionFlowNet compared to the fine-tune model without multitask learning. If blue is on top of the bar, our ActionFlowNet improves the accuracy, and vice versa.

4.4. Effects of Optical Flow

In this section, we study the effects of different optical flow models for action recognition.

We train our optical flow models on FlyingChairs or UCF101 and evaluate their accuracies on the Sintel dataset (similar to [5] that trains the model on FlyingChairs but tests on other datasets). The flow accuracy is evaluated by the end-point-error (EPE) measure, which calculates the averaged L_2 distance between the predicted and groundtruth flow vector at each pixel.

We investigate how the quality of the learned optical flow affects action recognition. Since optical flow in the multitask model is collaboratively learned with the recognition

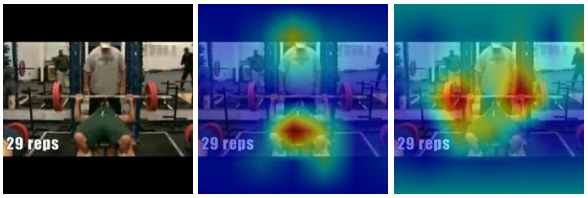

(a) Improved class: BenchPress (+29%)

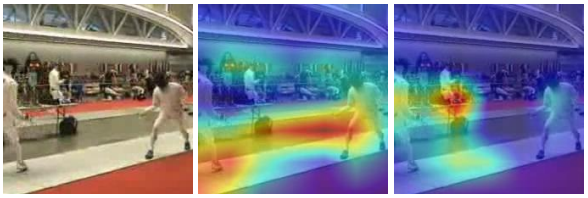

(b) Degraded class: Fencing (-15%)

Figure 5: Examples of improved/degraded classes and their importance visualization differences. In each group of three images, the left most image is the example frame. The center image is overlaid with importance obtained by the model trained without multitask. The right most image is overlaid with importance obtained by the model trained with both appearance and flow.

task, the quality of optical flow in the multitask model does not directly affect recognition accuracy. Thus, we use our Stacked model learned with different datasets, fix the optical flow part and train the classification part in the network shown in Figure 2. We compare the end-point-error of different optical flow learners and the corresponding classification accuracy in Table 3.

With accurate groundtruth labels, the FlowNet trained with FlyingChairs dataset performs quantitatively better than the one trained with UCF101 data: 9.12 vs 11.84 in EPE. The flow accuracies from our models are well behind the state of the art optical flow algorithms like EpicFlow, but they are competitive to a widely used optical flow algorithm like TV-L1 [30], which gives 10.46 EPE on Sintel dataset.

Action Recognition with Learned Flow. Surprisingly, even with lower end-point-error the Stacked model pretrained on FlyingChairs performs significantly worse than the one pretrained on UCF101 dataset (51.7% vs 69.6%), as shown in Table 3. Compared to the model directly taking high quality optical flow as input (77.7%), our models are still not as good as training directly on optical flow. We believe this is because the quality of learned optical flow is not high enough.

To understand how the learned optical flow affects action recognition, we qualitatively observe the optical flow performance in Figure 6. Even though the end-point error of the FlowNet pretrained on FlyingChairs is low, the estimated optical flow has lots of artifacts in the background

Method	EPE on Sintel	Classification Accuracy (%)
Stacked on FlyingChairs	**9.12**	51.7
Stacked on UCF101	11.84	**69.6**
ResNet on EpicFlow	6.29	77.7

Table 3: Comparison between End-Point-Error (EPE, lower is better) and the classification accuracy. Interestingly, better optical flow does not always result in better action recognition accuracy. Refer to the text for discussion.

and the recognition accuracy on top of that is correspondingly low. We believe the reason is that the FlyingChairs dataset mostly consists of large displacement flow. As a result, the FlyingChairs model performs badly on estimating small optical flow, which contributes less in the EPE metric when averaged over the whole dataset. This is in contrast to traditional optimization based optical flow algorithms that can predict small displacement well but having difficulties for large displacement.

In addition, traditional optical flow algorithms such as TV-L1 and EpicFlow explicitly enforce smoothness and constancy. They are able to preserve object shape information when the flow displacements are small, which is important for action recognition. On the other hand, while our models perform comparably to traditional optical flow algorithms in terms of endpoint error, our model is not optimized for preserving flow smoothness. This shows that end-point-error of optical flow is not a good indicator of action classification performance, since shape preservation is not accounted in the metric.

We additionally trained the UCF101 FlowNet-Stack model in an end-to-end manner, where the pretrained FlowNet at the bottom is fine-tuned while training the CNN on top. We find that the model performs significantly worse, with 56.2% accuracy because of overfitting and also forgetting the previously learned knowledge - this reflects the difficulty of learning motion features directly from raw frames.

4.5. Visualizing Learned Filters

To better understand what is learned from the network, we visualize and compare the filters of the first convolutional layer of different models in Figure 7. Several interesting insights can be observed from the visualization.

First, the filters from our models are very different from the ImageNet pretrained network. The first layer filters of the ImageNet pretrained network mostly consist of color and edge detectors, which are relatively smooth. In contrast, the filters learned from FlowNet contains more high frequency patterns. We believe that unlike object classification in ImageNet, where object class irrelevant details can be discarded in the early layers, high frequency details are critical to obtaining accurate optical flow estimation.

(a) Frame with movements highlighted with red

(b) EpicFlow

(c) FlowNet on FlyingChairs

(d) FlowNet on UCF101

Figure 6: Qualitative comparison of flow outputs. It shows an example of small motion, where the maximum magnitude of displacement estimated from EpicFlow is only about 1.6px. FlowNet trained on FlyingChairs dataset fails to estimate small motion, since the FlyingChairs dataset consists of large displacement flow.

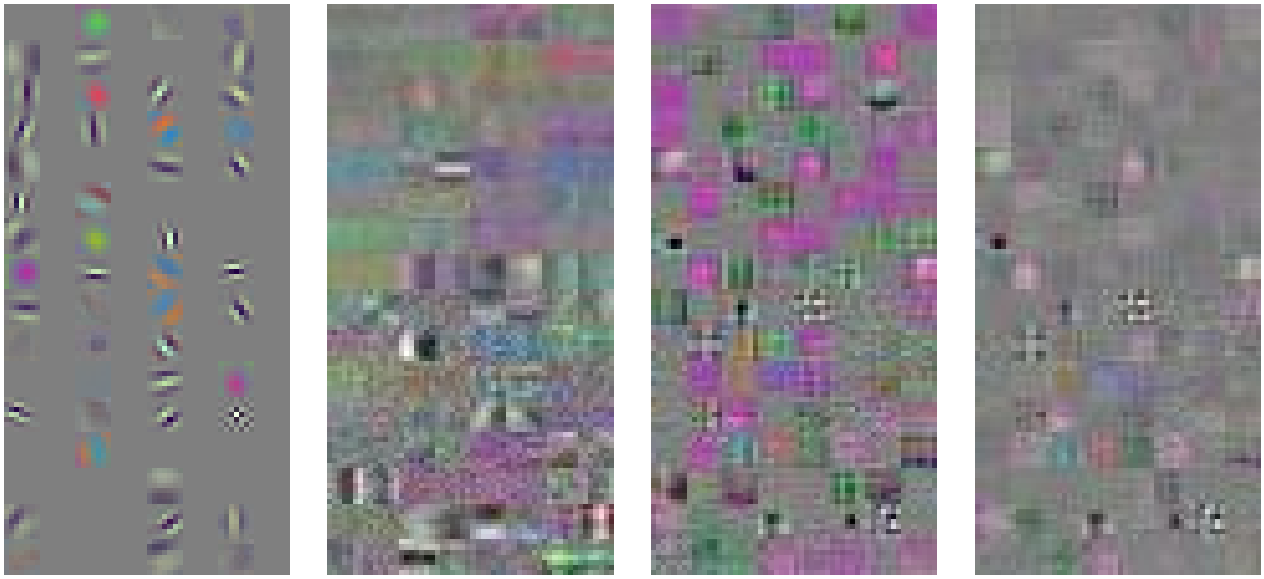

(a) ImageNet pretrained

(b) UCF101 from scratch

(c) FlowNet

(d) ActionFlowNet on UCF101

Figure 7: Visualization of our the learned filters. Every pair of horizontally adjacent patches corresponds to a single filter of the two input images; i.e. each row contains 4 filters. Refer to Section 4.5 for discussion.

Second, in our FlowNets (Figure 7c, 7d) the structures for the two input patches in each filter are generally different. This shows that our models are learning to capture motion, where the difference between two images is modeled, in addition to learning the appearance. Compared to the model that is trained from scratch, where two weight patches shares similar patterns, it shows that our flow-based models learn substantially different features.

While the filters learned from our FlowNet are quite different from the ImageNet pretrained network, our models perform reasonably well on action recognition, suggesting that Gabor like filters may not be sufficient for video classification. Instead, capturing motion requires different information from videos, which suggests that the current image based CNN architecture, while performing reasonably well on action recognition, may not be in the best form to learn motion representation.

5. Conclusion

We presented a multitask framework - *ActionFlowNet* for jointly estimating optical flow and video action classification. By using optical flow as supervision for classification, our model captures motion information while not requiring explicit optical flow computation as input. Our model significantly improves action recognition accuracy over the single task baselines with faster testing because additional optical flow computation is avoided.

Our work provide important insights on using optical flow to learn motion representation. Future work includes designing better model to exploit motion for learning video features. Different interesting extensions to incorporate multiple frames, including aggregating temporal information by feature pooling layers or LSTM [14], or apply 3D convolutions to hierarchically aggregated temporal information [21, 23], could be explored.

Acknowledgments

This research was supported in part by funds provided from the Office of Naval Research under grant N000141612713 entitled "Visual Common Sense Reasoning for Multi-agent Activity Prediction and Recognition".

References

[1] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, 2012. 4

[2] G. Cheng, Y. Wan, A. N. Saudagar, K. Namuduri, and B. P. Buckles. Advances in human action recognition: a survey. *arXiv preprint arXiv:1501.05964*, 2015. 1, 2

[3] A. Diba, A. M. Pazandeh, and L. V. Gool. Efficient two-stream motion and appearance 3d cnns for video classification. *arXiv preprint arXiv:1608.08851*, 2016. 2

[4] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016. 2

[5] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazırbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning Optical Flow with Convolutional Networks. In *ICCV*, 2015. 1, 2, 3, 4, 6

[6] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. *arXiv*, 2015. 2, 3, 4

[7] M. Jain, J. van Gemert, and C. Snoek. What do 15,000 object categories tell us about classifying and localizing actions? In *CVPR*, 2015. 6

[8] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale Video Classification with Convolutional Neural Networks. In *CVPR*, 2014. 1, 2

[9] D. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015. 4

[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*, 2012. 3, 4

[11] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *ICCV*, 2011. 4

[12] I. Misra, A. Shrivastava, A. Gupta, and M. Hebert. Cross-stitch Networks for Multi-task Learning. In *CVPR*, 2016. 4

[13] I. Misra, L. C. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016. 2, 5

[14] J. Y.-H. Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond Short Snippets: Deep Networks for Video Classification. In *CVPR*, 2015. 2, 8

[15] A. Rasmus, H. Valpola, M. Honkala, M. Berglund, and T. Raiko. Semi-Supervised Learning with Ladder Networks. In *NIPS*, 2015. 4

[16] E. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. EpicFlow: Edge-Preserving Interpolation of Correspondences for Optical Flow. In *CVPR*, 2015. 2, 4

[17] K. Simonyan and A. Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. In *NIPS*, 2014. 1, 2, 3, 5

[18] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2014. 3

[19] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A Dataset of 101 Human Action Classes From Videos in The Wild. In *CRCV-TR-12-01*, 2012. 4

[20] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going Deeper with Convolutions. In *CVPR*, 2015. 4

[21] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning Spatiotemporal Features with 3D Convolutional Networks. In *ICCV*, 2015. 1, 2, 8

[22] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Deep End2End Voxel2Voxel Prediction. In *CVPRW DeepVision Workshop*, 2016. 4

[23] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *arXiv preprint arXiv:1604.04494*, 2016. 8

[24] C. Vondrick, H. Pirsiavash, and A. Torralba. Generating videos with scene dynamics. In *NIPS*, 2016. 2, 5

[25] H. Wang, A. Klaser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *CVPR*, 2011. 1

[26] H. Wang, M. M. Ullah, A. Kl"aser, I. Laptev, and C. Schmid. Evaluation of local spatio-temporal features for action recognition. In *BMVC*, 2009. 1

[27] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 2

[28] X. Wang, A. Farhadi, and A. Gupta. Actions ˜ transformations. In *CVPR*, 2016. 2

[29] X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *ICCV*, 2015. 2

[30] C. Zach, T. Pock, and H. Bischof. A Duality Based Approach for Realtime TV-L1 Optical Flow. In *Pattern Recognition*. 2007. 7

[31] M. D. Zeiler and R. Fergus. Visualizing and Understanding Convolutional Networks. In *ECCV*, 2014. 5